Joint Impedance and Facies Inversion – Seismic inversion redefined

Michael Kemper[1*] and James Gunning[2] present a new seismic inversion system, which aims to overcome shortcomings in present-day simultaneous inversion tools by using joint impedance and facies inversion technology.

In this paper we will first review the industry-standard simultaneous inversion method (which derives continuous impedances) and subsequently identify some pitfalls. We will then introduce our new Joint Impedance and Facies Inversion technology (which we call Ji-Fi for short in this paper), which overcomes these pitfalls by recasting the seismic inverse problem as mixed discrete/continuous. Having so captured the correct physics, we apply this first on a wedge model, followed by a case study, before drawing some conclusions.

Note that in this paper, it is assumed that the seismic to be inverted is an ensemble of true amplitude partial angle stacks with corresponding wavelets derived from well ties.

Simultaneous inversion to date

Deterministic method

We start in this case with the Fatti (1994) equation (other approximations to the Zoeppritz (1919) equation such as Aki and Richards (1980) or Bortfeld (1961) are equally valid starting points):

$$R_{pp}(\theta) = aR_{AI} + bR_{SI} + cR_{\rho} \quad (1)$$

Where:

- $R_{AI} = \Delta Vp/(2Vp) + \Delta\rho/(2\rho)$
- $R_{SI} = \Delta Vs/(2Vs) + \Delta\rho/(2\rho)$
- $R_{\rho} = \Delta\rho/(2\rho)$
- $a = 1 + \tan^2\theta$
- $b = -8\,(Vs/Vp)^2 \sin^2\theta$
- $c = 4\,(Vs/Vp)^2 \sin^2\theta - \tan^2\theta$

We can turn reflectivity $R_{pp}(\theta)$ in (1) into synthetic seismic $S_{pp}(\theta)$ by convolving with wavelet $W(\theta)$:

$$S_{pp}(\theta) = a\,W(\theta)R_{AI} + b\,W(\theta)R_{SI} + c\,W(\theta)R_{\rho} \quad (2)$$

Subsequently we can use the small contrast approximation $R_{AI} = (AI2 - AI1) / (AI2 + AI1) = \frac{1}{2}\,\Delta AI / AI \approx \frac{1}{2}\,\Delta Ln(AI)$, same for R_{SI} and R_{ρ}, to rewrite (2) as

$$S_{pp}(\theta) = a/2\,W(\theta)\,\Delta Ln(AI) + b/2\,W(\theta)\,\Delta Ln(SI) + c/2\,W(\theta)\,\Delta Ln(\rho) \quad (3)$$

We know that if we proceed on the basis of (3), we obtain results that are not credible when compared to the impedances at the wells, as the seismic inverse problem is very underdetermined. So a first element of regularization is often introduced at this stage, often taking the form of a 'trick'. For instance, Ln(SI) and Ln(ρ) are commonly linearized with respect to Ln(AI):

$$Ln(SI) = \alpha_{SI}\,Ln(AI) + \beta_{SI} + \delta Ln(SI) \quad (4a)$$

$$Ln(\rho) = \alpha_{\rho}\,Ln(AI) + \beta_{\rho} + \delta Ln(\rho) \quad (4b)$$

Where:

- α_{SI}, β_{SI} are slope and intercept of the SI vs. AI linearization
- α_{ρ}, β_{ρ} are slope and intercept of the ρ vs. AI linearization
- $\delta Ln(SI)$ is the deviation away from the SI vs. AI linearization (as Ln(SI) is not quite linearly related to Ln(AI) – AI and SI contain independent information)
- $\delta Ln(\rho)$ is the deviation away from the ρ vs. AI linearization (same remark as above)

Inserting (4a) and (4b) into (3) leads to

$$S_{pp}(\theta) = a'/2\,W(\theta)\,\Delta Ln(AI) + b/2\,W(\theta)\,\Delta\delta Ln(SI) + c/2\,W(\theta)\,\Delta\delta Ln(\rho) \quad (5)$$

Where:

- $a' = a + \alpha_{SI}\,b + \alpha_{\rho}\,c$

Equation (5) indicated that we invert to Ln(AI), δLn(SI) and δLn(ρ), i.e. for Ln(SI) and Ln(ρ) we invert for the *deviation* from the linearizations expressed in (4a, 4b).

Because of the convolution by W(θ), equation (5) cannot be applied sample-by-sample. Instead it is used trace-by-trace (this can be extended to multi-trace; not discussed here). Therefore Ln(**AI**), δLn(**SI**) and δLn(ρ) can be seen as column vectors which we stack on top of one another to

[1] *Ikon Science Ltd.*
[2] *CSIRO.*
[*] *Corresponding Author, E-mail: mkemper@ikonscience.com*

form a block column vector called Ln(**Z**). The difference operation Δ, applied to each of Ln(**AI**), δLn(**SI**) and δLn(ρ), can be expressed as an almost diagonal matrix **D**. And lastly the convolution can be expressed by a banded matrix **W** (typically a different **W** matrix per partial angle stack). The product of **W** and **D** we can call system matrix **J**, into which the scaling parameters a, b and c (and the ½ factor) can be subsumed. Subsequently a block system matrix **J** is formed from system matrices **J**, so that equation (5) in block matrix form reduces to

$$\mathbf{S} = \mathbf{J} \bullet \mathrm{Ln}(\mathbf{Z}) \qquad (6)$$

Where:

- **S** is also a block column vector of the near, mid, far etc seismic traces stacked on top of one another.

Equation (6) is the essence of many *continuous* simultaneous inversion methodologies (although there are other implementations). To obtain Ln(**Z**), from which **Z** can be obtained by exponentiation, you either need to perform matrix inversion (quite expensive as the system matrix **J** can be large; the fact that **J** is typically singular exacerbates this) or you plug in a starting model of **Z** (which we label **Zo**), take the natural logarithm, multiply by **J** from the left, and for each partial angle stack incidence angle θ compare the synthetic seismic $\mathbf{S}_{pp}(\theta)$ so obtained with the real seismic $\mathbf{S}_{real}(\theta)$. Then use some optimization apparatus (e.g. the conjugate gradient method, or least squares optimization) to iteratively change **Z** until the difference between $\mathbf{S}_{pp}(\theta)$ and $\mathbf{S}_{real}(\theta)$ is minimized for all partial angle stack incidence angles θ.

So the optimization consists of minimizing an objective function, typically $\|\mathbf{S}_{real} - \mathbf{X} \bullet \mathbf{Ln}(\mathbf{Z})\|^2$, by judiciously changing Ln(**Z**). Even though in (4a) and (4b) we introduced a form of regularization, using this objective function is still unstable in practice, so we need to add a second level of regularization in the form of another, second term in the objective function to ensure that the impedances do not drift away too much from the initial impedances **Zo**. The total objective function to be minimized is then something like $\|\mathbf{S}_{real} - \mathbf{X} \bullet \mathrm{Ln}(\mathbf{Z})\|^2 + \mu\|\mathrm{Ln}(\mathbf{Z}) - \mathrm{Ln}(\mathbf{Zo})\|^2$, where μ, the so-called model weight, should be as small as possible to ensure the inversion is driven mostly by the data (i.e. the seismic $\mathbf{S}_{real}$), and as little as possible by the initial model (**Zo**). Note that in this case the objective function is quadratic, which makes optimization straightforward. Sometimes L_1 norms are used for the second term, making optimization somewhat more difficult.

Statistical method

Clearly, simultaneous inversion can be seen as a statistical problem, given the noise component of the seismic signal, uncertainty in rock properties etc. Realising this, the seismic inversion can be recast as a Bayesian problem, in which the prior information, if well chosen, will provide adequate regularization. Bayes's Theorem in this case can be written as:

$$\pi(\mathbf{Z}|\mathbf{S}_{real}) \approx L(\mathbf{S}_{real}|\mathbf{Z})\ p(\mathbf{Z}) \qquad (7)$$

Where:

- $\pi(\mathbf{Z}|\mathbf{S}_{real})$ is the posterior distribution
- $L(\mathbf{S}_{real}|\mathbf{Z})$ the likelihood function
- $p(\mathbf{Z})$ the prior distribution.

It is customary (Buland and Omre, 2003) to represent the distributions as being multi-normal; this is often very reasonable, and makes the mathematics more tractable.

The prior distribution p(**Z**) can be obtained from a linear depth trend of, say, AI and cross-plots between AI vs. SI and AI vs. ρ (all typically derived from well data but could come from simple 1D basin modelling or from an analogue database, and all with an assessment of uncertainty). These trend fits can be expressed as a multi-normal prior distribution p(**Z**) of form

$$p(\mathbf{Z}) \sim \exp\{-\tfrac{1}{2}(\mathbf{Z} - \mathbf{Zo})^T\ \mathbf{C}_p^{-1}\ (\mathbf{Z} - \mathbf{Zo})\} / |\mathbf{C}_p|^{\frac{1}{2}} \qquad (8)$$

Where:

- $\mathbf{C}_p$ is the covariance matrix describing the variance of and the correlation between the impedances

The likelihood function $L(\mathbf{S}_{real}|\mathbf{Z})$ can be expressed as

$$L(\mathbf{S}_{real}|\mathbf{Z}) \sim \exp\{-\tfrac{1}{2}(\mathbf{S}_{real} - F(\mathbf{Z}))^T\ \mathbf{C}_d^{-1}\ (\mathbf{S}_{real} - F(\mathbf{Z}))\} / |\mathbf{C}_d|^{\frac{1}{2}} \qquad (9)$$

Where:

- F(**Z**) is the function to derive synthetic seismic from the impedances **Z**, as described above under 'Deterministic method'
- $\mathbf{C}_d$ is the covariance matrix representing the seismic noise.

The (un-scaled) posterior distribution can be derived using (7), from which we can derive the maximum a-posteriori (MAP) model of **Z** or we can use McMC sampling to obtain marginal distributions of interest (Gunning and Glinsky, 2004).

Pitfalls of conventional simultaneous inversion

Even though Simultaneous Inversion has been the 'workhorse' seismic inversion tool for many years, there are some shortcomings that we will point out in this section:

Inadequate regularization

Equation (4a), used in deterministic simultaneous inversion, suggests that the Ln(SI) / Ln(AI) relationship can be expressed as one straight line (same for (4b)). However, from Figure 1, where Ln(SI) for a typical dataset is cross-plotted as a function of Ln(AI), coloured by the volume of shale,

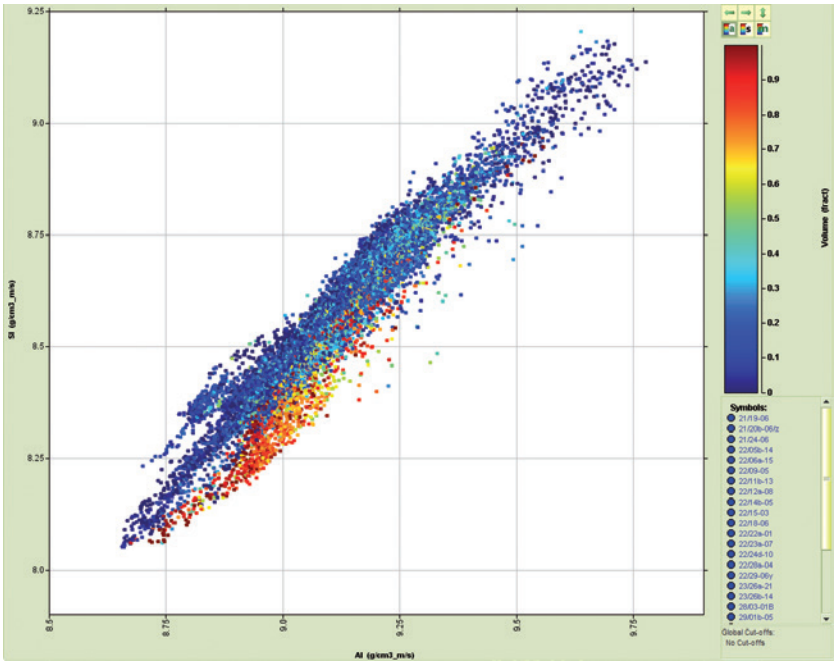

Figure 1 *Ln(SI) / Ln(AI) cross-plot, colour coded by Vshale.*

it should be clear to any rock physicist that three lines are required, two through the blue sand prone data cloud (water bearing and hydrocarbon bearing sands), and one through the orange shale prone data points. One line clearly does not suffice. So ideally we would like to see Rock Physics relationships *per facies*.

Regularization in the Bayesian form of simultaneous inversion is provided more elegantly and (in many cases) more accurately compared to the deterministic approach by use of the prior information – equation (8). However, these priors are derived as linear trends of AI (in case of Fatti) with TWT, and it is unlikely that this represents the AI low frequency behaviour adequately. In fact, such a linear depth trend may only represent frequencies from 0 to, say, 2 Hz, and if the seismic to be inverted contains frequencies from, say, 5 Hz onwards, there will be a frequency deficiency in the inverted impedances from 2 to 5 Hz in this example. If the phenomena to be studied are outside this frequency range, then this may not be an issue, but it is clearly worth checking beforehand. Ideally, of course, we would like to use prior regularization that is not frequency deficient.

Inadequate low frequency background model (LFBM) building

Although more advanced ways of building LFBM's have been published (Sams and Saussus, 2013; Douma and Naeini, 2014), the most popular route is to interpolate low-pass filtered impedance profiles derived at the wells. The challenge is that if you have, say, a high Net-to-Gross (NtG) well in the West and a low NtG well in the East, how would you interpolate a credible LFBM impedance profile at a location in between? You would really need to know NtG at that location, and of course you do not know that beforehand (determining the NtG may be one of the reasons why you want to perform seismic inversion in the first place). Ideally, we would like to have a way of interpolating well derived impedance information in a manner that does not depend on NtG. The way to do so is to interpolate the impedance data *per facies*. In, for instance, a binary sand/shale sequence, the sand and shale depth trends individually are quite well constrained from well data, and these can then be readily interpolated to unknown locations. Importantly, knowledge of NtG is now not a requirement.

Physics not properly captured

In Figure 2, we make an AI 'truth' model in black (also SI and Rho, not shown), we synthesize the seismic (only the zero incidence synthetic is shown at the right) and then perform a simultaneous inversion back to AI in red (and SI and Rho, not shown), using the dashed line as the AI LFBM. Note that the biggest loops in the synthetic are where we switch between the facies (think of this problem as for instance a Shale/Sand/Shale sequence), and also note that simultaneous inversion only gives credible results near these facies switches; away from these facies switches, the AI result trends back to the LFBM, as indicated by the arrows.

Clearly the facies, which can be seen as discrete quantities, have a major control on the seismic expression, but these are not inverted for in standard simultaneous inversion. So ideally we would like to invert for these (discrete) facies also, as well as of course for the (continuous) impedance *per facies*. Only if we invert for both can the physics be captured properly.

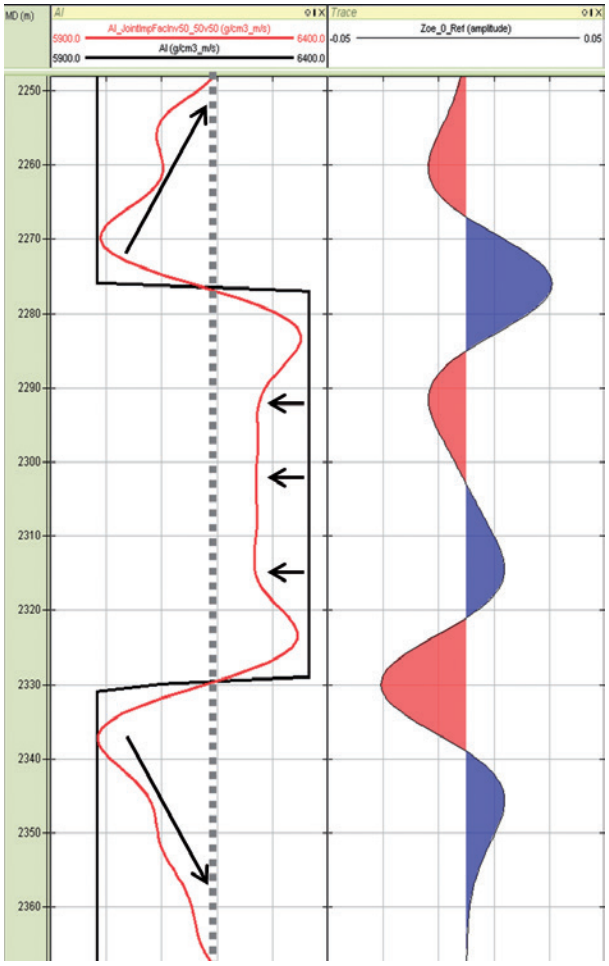

Figure 2 *A synthetic AI model (black) with corresponding synthetic and simultaneously inverted AI profile (red). See text for details.*

Joint categorical/continuous simultaneous inversion

In the previous section we have learnt that, ideally, to capture the physics of the seismic inverse problem properly, we invert to facies and impedances *per facies*, and that *per facies* Rock Physics Models and LFBM's should be used.

To do so is mathematically demanding, as it cannot be written down in closed form (such as (6) above), nor can it be solved using standard optimization apparatus. We use an iterative method where we first invert the seismic for facies (given a bland impedance model), then given these facies we invert the seismic for impedances, then given these impedances we invert the seismic for facies and so forth, until convergence. This optimization is a form of Expectation-Maximization (McLaughlan, 2000): in the Expectation step we invert the seismic for facies given the impedances, and in the Maximization step we invert the seismic for impedances given the facies.

Again using Bayes' theorem, the posterior distribution equivalent of (7) is now more complex

$$\pi(\mathbf{Z},F|\mathbf{S}_{real}) \approx L(\mathbf{S}_{real}|\mathbf{Z})\ p(\mathbf{Z}|F)\ p(F) \qquad (10)$$

The likelihood function $L(\mathbf{S}_{real}|\mathbf{Z})$ is the same as (9), and the prior distribution $p(\mathbf{Z}|F)$ is very similar to (8), the only difference being that the prior mean $\mathbf{Zo}$ and covariance matrix $\mathbf{C}_p$ now depends on facies F:

$$p(\mathbf{Z}|F) = \exp\{-½(\mathbf{Z} - \mathbf{Zo}(F))^T\ \mathbf{C}_p(F)^{-1}\ (\mathbf{Z} - \mathbf{Zo}(F))\}\ /\ |\mathbf{C}_p(F)|^{½} \qquad (11)$$

This *facies-dependent* prior distribution can be obtained from depth trends of AI and cross-plots between AI vs. SI and AI vs. ρ (in case of Fatti), complete with an assessment of uncertainty: these are typically derived from well data but could come from simple 1D basin modelling or from an analogues dataset. The difference with section 1 is that now we develop these depth trends and cross-plots *per facies*. This is a crucial improvement.

That leaves p(F), the facies prior distribution. For this we use a discrete Markov Random Field. Expressed simply, any 3D seismic lattice consists of many pixels. We define a set of pixel-pairs connecting each pixel to its direct neighbours within in the inversion window. So each pixel belongs to six pixel-pairs (except at edges, corners). The probability of a configuration F_c of the whole lattice is then defined by the sum of potential energies over all the pixel-pairs in what is called a Gibbs distribution:

$$p(F) \approx \exp(-\Sigma V_p(F_p)) \qquad (12)$$

Where

- V_p represents the 'potential energy' of the set of facies F_p seen by each pixel-pair c.

For each clique, we can write the potential energy as $V_p = \beta\ I(F_{centre}, F_{neighbour})$ where the discrete indicator function I is 0 if facies F_{centre} and $F_{neighbour}$ are the same and is 1 if they are not, and where β is a positive continuity parameter. So we can rewrite (12) as ...

$$p(F) \approx \exp(-\Sigma\ \beta\ I(F_1,F_2)) \qquad (13)$$

... and thus we penalize (reduce) the probability p(F) any time two neighbouring facies are different.

We have implemented different β›s for horizontal (inline and crossline direction) and vertical continuity, as geologically horizontal continuity is typically larger than in the vertical direction. However, as geology is seldom perfectly horizontal, we use the concept of chronostratigraphic age (Gawith et al., 2013) to determine neighbours of the same age, which may not be simply the neighbouring pixel at the same time index.

Note that the Markov Random Field also allows us to prohibit illegal facies combinations, i.e. we can ensure that we never obtain water-bearing sand on top of gas-bearing sand for instance. Also, note that this technique is different

to geostatistical inversion to facies and impedances or rock properties (see e.g. Sams et al., 2011) in that no variography is used.

Again, the un-scaled posterior distribution can be derived using (10), from which we can derive the maximum a-posteriori (MAP) model of **Z** and F, or we can use McMC sampling to obtain marginal distributions of interest.

Note: Below Figure 1 we indicated that the regularization of the Bayesian form of simultaneous inversion may lead to a frequency deficiency. Ji-Fi is clearly an expansion of the Bayesian form of simultaneous inversion (compare equation (10) to (7)), so could Ji-Fi suffer from this frequency deficiency also? The big difference between the two methods is that in the Bayesian form of simultaneous inversion you specify one linear trend of AI with Depth, whereas in Ji-Fi you specify two or more such trends (as in Ji-Fi you always invert to two or more facies). This means that within the inversion window the impedance results switch between the per facies trends, and this switching typically ensures that enough frequencies are injected to cover the frequencies from 0 Hz to the start of the seismic bandwidth.

Wedge model

We apply both conventional simultaneous inversion and Ji-Fi to a sand wedge model encased in a shale (Figure 3). The one LFBM used in simultaneous inversion has a Vp exactly in between the Vp of shale and the Vp of sand (i.e. is green against e.g. the colour bar of Figure 3 Top Left). The two LFBM's used in Ji-Fi have the Vp of shale and the Vp of sand respectively (i.e. dark red and dark blue respectively). Clearly, Ji-Fi gives a superior inversion result, but of course this is a conceptual model, so does not prove that Ji-Fi performs better on real data. We'll examine this in the next section.

Case study

We apply both conventional simultaneous inversion and Ji-Fi to a Triassic oil and gas field offshore Western Australia. We have defined 4 facies to invert for (Shale, Wtr Sand, Oil Sand and Gas Sand) from five wells that penetrate this field. Near, mid and far partial angle stack cubes are available, properly pre-conditioned for seismic inversion (not shown). Representative wavelets for the three partial angle stacks are estimated from well ties using the White (1980) method, which includes a seismic noise estimate. To compare the Ji-Fi facies result against simultaneous inversion (which does not give a direct facies estimate), we have performed Bayesian classification to facies on the simultaneous inversion impedances as a post-processing step.

In Figure 4 we show the one AI LFBM required for simultaneous inversion (left), as well as two of the four AI LFBM's required for Ji-Fi. Even though multiple LFBM's are required for Ji-Fi and only one for simultaneous inversion, it should be clear that deriving the multiple *per facies* Ji-Fi LFBMs is an easy-to-accomplish exercise, and that these LFBM's do not suffer from the interpolation halos you see on the one simultaneous inversion LFBM (which was derived by well log interpolation).

In Figure 5 we show (top left) a horizon slice through the facies cube obtained from simultaneous inversion followed by Bayesian classification, and bottom left the equivalent horizon slice through the Ji-Fi facies cube. Seismic Net Pay determination can be a complex endeavour (see Connolly, 2007; Connolly and Kemper, 2007), but once you have facies cubes, determining the net sand is equivalent to simply counting, per trace (over the inversion window), the number of Wtr Sand, Oil Sand and Gas Sand occurrences, and this is displayed to the right of this figure (top: simultaneous inversion followed by Bayesian classification; bottom: Ji-Fi).

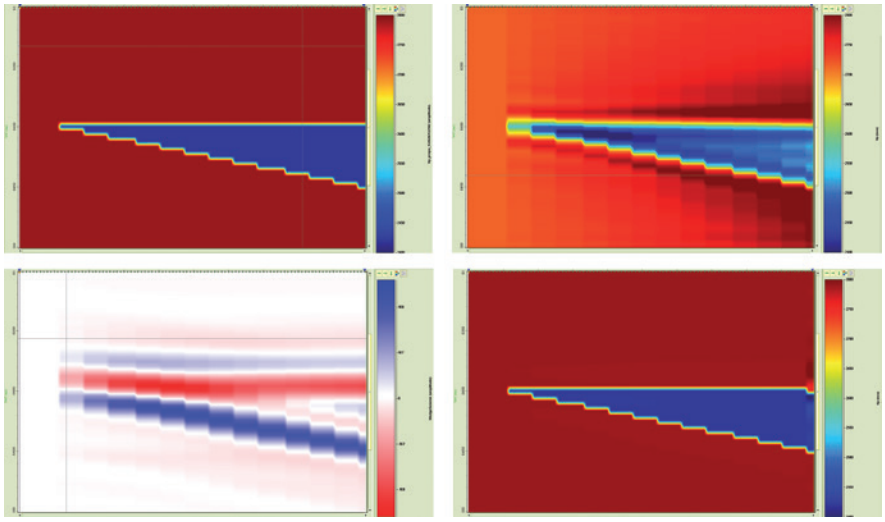

Figure 3 *Top Left: a 'truth model' for Vp (Vs and ρ not shown in this collage). Bottom Left: the corresponding zero incidence seismic (mid and far stacks not shown). Top Right: the Vp result of simultaneously inverting the synthetic seismic. Bottom Right: the Vp result of Ji-Fi inverting the synthetic seismic (Ji-Fi additionally gives a facies image – not shown).*

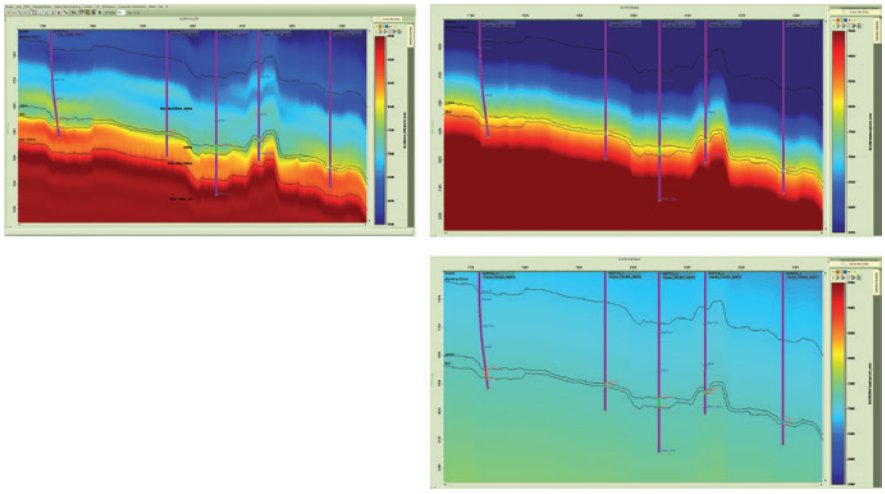

Figure 4 *Left: the AI LFBM used for simultaneous inversion. Right: two of the four AI LFBM's used for Ji-Fi (Top: Shale; Bottom: Wtr Sand; Oil Sand and Gas Sand not shown as they are very similar to WtrSand (shifted to lower AI values).*

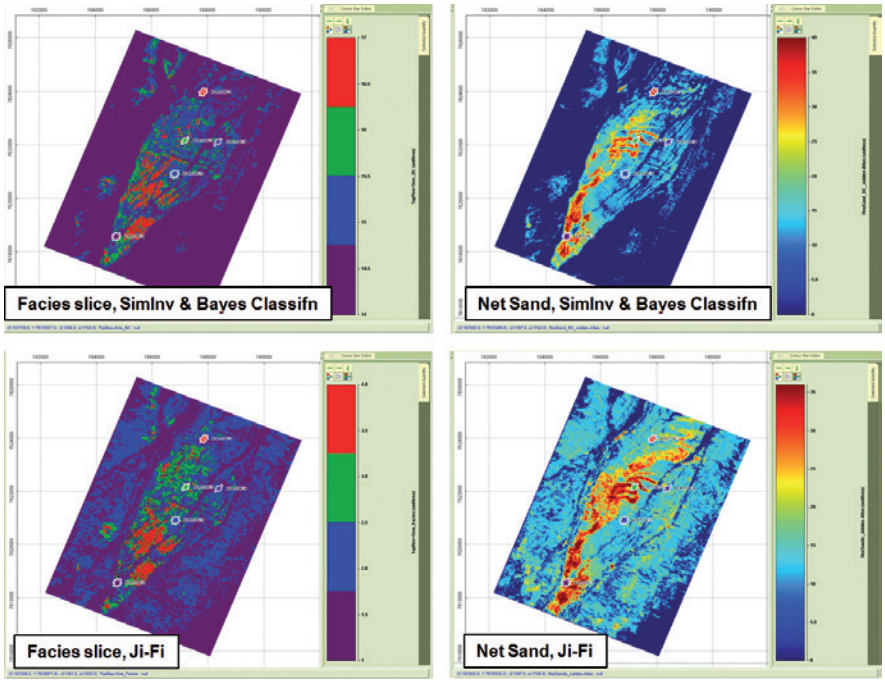

Figure 5 *Simultaneous inversion and Ji-Fi results on a Triassic oil and gas field offshore Western Australia. See text for details.*

In our opinion Ji-Fi gives a better inversion result, in that only Ji-Fi:

- Shows a proper match as regards hydrocarbon content to all wells,
- Images the channel as a nice, continuous feature,
- Finds water-bearing sands off structure (where you would expect them).

Note that in a future companion paper we shall present more case studies.

Conclusions

In-depth analysis of the present industry standard continuous simultaneous inversion method highlighted some shortcomings, which require the following remedies:

1. We need to invert to facies and impedances *per facies* (to capture the physics correctly)
2. We need to use *per facies* LFBMs

In this paper we have introduced a new joint discrete/continuous simultaneous inversion method which implements these two remedies, and which does not show unsightly artefacts of standard methods such as impedances tending to a non-geological value away from seismic energy, frequency deficiencies, or interpolation halos.

A particularly pleasing aspect of Ji-Fi is that well data is used in a 'soft' way only, making QC at the wells blind well QC by design.

Comparisons to a wedge model and to case studies (of which only one is shown in this paper; more will be

presented in a future companion paper) indicates that Ji-Fi gives an improved inversion result compared to simultaneous inversion.

Acknowledgements

I would like to thank Tullow Oil for sponsoring this project. I also thank my colleagues Ana Somoza and Kester Waters for assisting with some of the figures.

References

Aki, K. and Richards, P.G. [1980] *Quantitative seismology: Theory and Methods*. W.H. Freeman and Co., New York, 932.

Bortfeld, R. [1961] Approximation to the reflection and transmission coefficients on plane longitudinal and transverse waves. *Geophysical Prospecting*, **9**, 485–503.

Buland, A. and Omre, H. [2003] Bayesian linearized AVO inversion, *Geophysics*, **68**, 185–198.

Connolly, P.A. [2007] A simple, robust algorithm for seismic net pay estimation. *The Leading Edge*, **26**(10), 1278-1282.

Connolly, P.A. and Kemper, M.A.C [2007] Statistical uncertainty of seismic net pay estimations. *The Leading Edge*, **26**(10), 1284-1289.

Douma, J. and Naeini, E. [2014] Application of image-guided interpolation for building LFBM prior to inversion. *76th EAGE Conference & Exhibition,* Extended Abstracts.

Fatti, J.L., Smith, G.C., Vail, P.J., Straus, P.J. and Levitt, P.R. [1994] Detection of gas in sandstone reservoir using AVO analysis: a 3D seismic case history using the Geostack technique, *Geophysics*, **59**, 1362-1376.

Gawith, D.E., Gutteridge, P. and Kemper, M. [2013] The Advantages of Gridless Geological Modelling for Reservoir Surveillance. *83rd SEG Annual Meeting,* Expanded Abstracts.

Gunning, J. and Glinsky, M.E. [2004] Delivery: an open–source model–based Bayesian seismic inversion program. *Computers & Geosciences*, **30**, 619–636.

McLaughlan, G. J. [2000] *Finite Mixture Models*. Wiley.

Sams, M.S., Millar, I., Satriawan, W., Saussus, D. and Bhattacharyya, S. [2011] Integration of geology and geophysics through geostatistical inversion: a case study. *First Break*, **29**(8), 1–10.

Sams, M. and Saussus, D. [2013] Practical implications of low frequency model selection on quantitative interpretation results, *83rd SEG Annual Meeting,* Expanded Abstracts.

White, R.E. [1980] Partial coherence matching of synthetic seismograms with seismic traces. *Geophysical Prospecting*, **28**(3), 333-358.

Zoeppritz, K. [1919] Erdbebenwellen VIII B. On the reflection and penetration of seismic waves through unstable layers. *Goettinger Nachrichten*, **1**, 66-84.

DEEPWATER SUCCESS. PLAN ON IT.

KNOW BEFORE YOU GO

In deep water, there's no room for guessing. While the rewards of ultra-deep reservoirs are huge, so are the risks. But you can plan for success with Ikon. We integrate our geopressure and geomechanics expertise with Quantitative Interpretation based on rock physics to provide a more accurate understanding of the subsurface, so you know where to drill and what to avoid. All backed by experts with a proven track record around the world. Take the pressure off your shoulders. Know before you go.

www.ikonscience.com

ikon Science

The Present And Future Of GeoPrediction